




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 13/16	A1	(11) International Publication Number: WO 98/15897 (43) International Publication Date: 16 April 1998 (16.04.98)
<p>(21) International Application Number: PCT/US97/18128</p> <p>(22) International Filing Date: 6 October 1997 (06.10.97)</p> <p>(30) Priority Data: 08/729,261 10 October 1996 (10.10.96) US</p> <p>(71) Applicant: HEWLETT-PACKARD COMPANY [US/US]; 3000 Hanover Street, Palo Alto, CA 94303 (US).</p> <p>(72) Inventor: WIGGERS, Hans, A.; 12110 Mellowood Drive, Saratoga, CA 95070 (US).</p> <p>(74) Agent: MAYER, Marc, R.; Hewlett-Packard Company, 1501 Page Mill Road, M/S 4U-10, Palo Alto, CA 94304-1126 (US).</p>	<p>(81) Designated States: JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: MEMORY SYSTEM AND DEVICE



(57) Abstract

This disclosure provides a memory system and device for synchronizing response across multiple memory devices (105), whether arranged serially upon a single data bus, in parallel across multiple data busses (107), or both. A memory controller (103) periodically configures the system (101) by separately placing each memory chip (105) into a configuration mode. While in this mode, the chip (105) is polled by the controller (103) along the corresponding data bus (107), and the chip (105) responds with a reply. The controller (103) uses this reply to compute elapsed time between polling and the reply. Using all of the chips (105), the controller (103) determines the maximum response time, in terms of elapsed clock cycles. Based on this maximum time, and the individual response times for each chip, the controller (103) then programs each chip (105) with a number which defines chip-based delay (149) for responses to data read operations. In this manner, successive data reads can be performed on successive clock cycles without awaiting prior completion of earlier data reads. In addition, in a multiple data bus system, the controller (103) is not delayed by having to wait for all simultaneous data reads across a wide bus.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

MEMORY SYSTEM AND DEVICE

This disclosure relates to a memory system and device and, more particularly, provides a method of synchronizing response from memory chips.

5

BACKGROUND

As computers and their central processing units ("CPUs") become capable of executing instructions more rapidly, this ability carries with it a need for increased memory size and speed, and also bus size. The need has
10 given rise to much design effort directed toward optimizing current and future memory device designs to provide quick memory response. Commonly-recognized current examples of memory devices include dynamic random access memories ("dRAMs"), read only memories ("ROMs") and synchronous
15 random access memories ("SRAMs"), as well as mechanical and optical devices, such as CD-ROMs.

In performing a typical data read operation, a memory controller (usually the CPU or a dedicated memory controller in larger systems) sends a read command to a
20 particular memory chip. This command is propagated to the chip along one or more lines of a command bus. When received by the particular chip, the command causes the chip to locate and direct an output from its internal memory array onto a data bus, as a return data signal
25 intended for the memory controller. The output then propagates along the data bus, which may or may not travel the same route as the command bus. In the example just given, there are four sources of time delay, including the propagation time of a read command from the controller to
30 the chip, the time required for the chip to power its

internal registers and channel the proper output onto the data bus, and the time required for propagation of the output back to the controller. The fourth source of time delay, present in systems which have split data busses (e.g., with separate, parallel data busses carrying bit-groups of different significance), is controller delay caused by slow retrieval along one of the parallel data busses. Typically, design efforts have focussed only on reducing the second of these times, that is, on improving internal routing and processing of instructions within memory chips.

The design efforts mentioned above, however, while continually providing more responsive memory devices, do not enable synchronization across multiple busses, nor do they eliminate the possibility of simultaneous bus contention among multiple memory chips, a problem explained with reference to FIGS. 1 and 2.

Bus Contention.

FIG. 1 illustrates a hypothetical memory system which is accessed by a CPU. The memory system includes a memory controller 23, a command bus 25, a data bus 27, and two memory devices 29 and 31 which are, in the example of FIG. 1, RAM chips. The command bus includes a system clock signal 33, issued by the controller, a data read signal 35, and an address bus 37. Each of the RAM chips 29 and 31 are presumed to be at different distances from the controller 23 from the standpoint of bus wiring, such that commands take a slightly different amount of time to reach each chip. FIG. 2 provides timing diagrams for purposes of explaining the problem of simultaneous bus contention.

It is desired for the memory system 21 to operate as quickly as possible, so the controller 23 will issue three data read commands on consecutive clock cycles; the first and third of these commands are directed to each of the aforementioned RAM chips 29 and 31, and the second command is unimportant to the present discussion and, so, it is indicated by a dash (-). It is presumed that the system clock has a frequency of 250 megahertz, so the width of each square pulse of FIG. 2A has an associated "high" time of 2 nanoseconds. FIG. 2B illustrates the issuance of two data read operations 39 and 41 by the memory controller, respectively labeled "X" and "Y," which then propagate along the command bus toward their intended RAM chip destinations.

FIG. 2C indicates timing of command "X" response by a first one 29 of the RAM chips, whereas FIG. 2D indicates timing of command "Y" response by the second one 31 of the RAM chips. The first chip 29, as indicated by FIG. 2C, receives its command "X" 8 clock cycles after it has been propagated, and beginning on the ninth clock cycle, takes 20 nanoseconds to retrieve and place its output "D-X" onto the data bus. By contrast, the second chip 31, as indicated by FIG. 2D receives its command "Y" only 7 clock cycles after it has been propagated, and it also takes 20 nanoseconds to place its output "D-Y" onto the data bus. If the output for each RAM chip 29 and 31 takes the same amount of propagation time to return to the controller as was the case for the commands to originally reach the RAM chips, then both outputs "D-X" and "D-Y" will arrive at the controller 23 at the same time, as indicated by a hatched block 43 of FIG. 2E. It will be noted that the address bus and data bus have been indicated, in this example, to have similar propagation times associated with them, though this result will not necessarily be the case in memory system

designs in which the command bus and data bus are separately routed. Further, while only two memory devices are indicated above, it will readily be seen that the problem of bus contention is particularly complicated as the number of memory devices is increased, and as CPUs become more and more efficient, operating at multi-hundred megahertz frequencies and greater.

There exists a definite need for a memory system which permits simultaneous response along multiple, parallel data busses. Also, there exists a definite need for a system that avoids simultaneous bus contention between multiple memory devices along one data bus. Preferably, such a system would address situations where propagation times for reaching individual chips may differ by an entire clock cycle or greater. By avoiding multiple device bus contention, it is hoped that memory controllers, such as CPUs and dedicated memory controllers, can perform data read and write operations on consecutive clock cycles, thereby keeping pace with CPU design, and the tendency toward increased computer speed and memory capacity. The present invention solves these needs and provides further, related advantages.

SUMMARY OF THE INVENTION

The present invention provides a memory system and device that fulfills the needs mentioned above; it provides a memory system where individual memory devices (such as micro-chips) can be delayed in operation to provide data outputs to a controller, such as a CPU or memory controller, a predetermined time after a request for data is made. Consequently, the CPU or memory controller does not have to wait until a first data read is completed prior

to beginning a second data read. Also, the system facilitates use of a scalable data bus, since memory chips feeding parallel 8-bit or other size data busses can be delayed to provide simultaneous response to a CPU,
5 notwithstanding differences in wiring and bus routing. As can be seen therefore, the present invention provides a memory system with substantial utility.

A first form of the invention provides a memory system having a controller and at least two memory devices,
10 with at least one data bus coupling the controller with each memory device. The controller sends a command to both of the memory devices, which provide a reply back to the controller. When sending a command, the controller begins timing, and awaits receipt of the reply, when the
15 controller is able to determine an elapsed time. Once this is done for both memory devices, the controller determines a maximum travel time and adjusts the response time of at least one of the two memory devices.

In this manner, a device built according to the
20 present invention can synchronize response for memory devices (1) across parallel data busses, for example, across four individual 16-bit data busses that make up a 64-bit data bus, to ensure simultaneous response, (2) serially, along one data bus, to avoid bus contention
25 between multiple devices along the same data bus, such that data reads can be performed on successive clock cycles, or (3) both of the above.

To this effect, a second form of the invention provides a memory device having a buffer for delaying
30 operations of the memory device. According to this form of the invention, the memory device receives a system clock and has stored in the buffer a number of clock cycles by

which it is to delay the device's output. When a data read command is then received, the device can first begin by counting until the number indicated by the buffer is reached, and then commence data fetch operations.

- 5 Alternatively, the device can immediately fetch data from an internal memory storage area and delay latching its output onto the corresponding data bus by an amount responsive to the buffer's contents.

10 In more particular aspects of the invention, the memory device can include a configuration mode which the device is directed to enter by a system controller. The controller preferably writes a configuration command to a mode register of the memory device, which causes the device to monitor the data bus connecting the device to the
15 controller; the device monitors the bus to detect both an input strobe and a predetermined data word from the controller. When the memory device detects this data word, it latches a reply word onto the data bus together with an output strobe, preferably derived locally by the device
20 from the system clock. The controller uses the output strobe to notify it that incoming data on the data bus is valid, and consequently may be used to determine elapsed time. From this elapsed time, the controller determines the delay time suitable to the particular device, and loads
25 it into the device's buffer. Preferably, this result is performed by detecting a maximum of all response times, and by loading into each buffer a number of clock cycles which delays latching of the corresponding chip's output to synchronize response to the maximum time.

30 In still more particular features of the invention, the device can perform correction to the buffer contents, by accepting a write command from the controller, by counting a difference between time of arrival of the write

command itself and write data, and by adjusting contents (the delay time) of the device's buffer, to thereby correct for differences in travel time between a command bus and a data bus.

5 Other forms of the invention are indicated below and in the claims, and include related systems and methods for performing memory synchronization related to the two forms of the invention discussed above.

10 The invention may be better understood by referring to the following detailed description, which should be read in conjunction with the accompanying drawings. The detailed description of a particular preferred embodiment, set out below to enable one to build and use one particular implementation of the invention, is not intended to limit
15 the enumerated claims, but to serve as a particular example thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of one hypothetical memory system used to illustrate the prior art; a broken
20 line separates a computer's central processing unit ("CPU"), which is above the line, and a memory system, located below the line.

FIG. 2 is a series of timing diagrams used to illustrate the problem of bus contention between multiple
25 memory devices, in the hypothetical prior art system of FIG. 1.

FIG. 3 is a block diagram of the preferred embodiment, which uses both parallel data busses ("baby-

busses," which together make up a wide system bus) and simultaneous fetch of words (of 64, 132 or greater bit length) and shared data bus use among multiple memory devices.

5 FIG. 4 is a block diagram of a single memory device of FIG. 3; preferably, the device is a dynamic random access memory ("DRAM").

 FIG. 5 is a series of timing diagrams corresponding to a portion of the system of FIG. 3; in particular, FIG. 5
10 illustrates synchronization of responses from two memory devices.

DETAILED DESCRIPTION

 The invention summarized above and defined by the enumerated claims may be better understood by referring to
15 the following detailed description, which should be read in conjunction with the accompanying drawings. This detailed description of a particular preferred embodiment, set out below to enable one to build and use one particular
implementation of the invention, is not intended to limit
20 the enumerated claims, but to serve as a particular example thereof. The particular example set out below is the specific implementation of a memory system, in particular, one that is used for to synchronize responses from multiple memory devices, whether arranged along a single data bus,
25 or in across parallel data busses, or both.

 In accordance with the present invention, the preferred embodiment is memory system 101, indicated in FIG. 3. A memory controller 103, which may be either a CPU or a dedicated memory controller, accesses numerous memory

chips 105 which are arranged both serially and in parallel. In other words, the system includes multiple, parallel data busses, preferably 16-bit data busses ("baby-busses") 107, with several memory devices on each baby-bus. This
5 construction lends itself to a scalable bus system, e.g., in which the system data bus is 64-bits or 132-bits wide, or greater, with the least significant bits of data being retrieved along one baby-bus 107-1, and the most significant bits of data being retrieved along another
10 baby-bus 107-2.

Implementing the features of the present invention, the controller synchronizes memory retrieval both amongst multiple chips 105 along each single baby-bus 107, as well as across the multiple baby-busses. Combined
15 synchronization both across parallel busses and across serially-arranged chips is not necessary to practice the invention as defined by the claims below, and synchronization of memory devices either serially or in parallel, alone, is within the spirit of the present
20 invention. Synchronization across multiple baby-busses 107 is performed in the preferred embodiment to ensure that the controller 103 is not delayed by a slow response along one of the baby-busses 107. Synchronization along each single baby-bus 107 is performed to reduce the possibility of bus
25 contention by multiple chips 105, enabling the controller to perform multiple data reads from different chips 105 on consecutive cycles notwithstanding differences in propagation times along each baby-bus 107.

The controller 103, if it is a dedicated memory
30 controller, receives memory requests along a master bus 109 which couples it to a system CPU (not illustrated). These requests are managed by the controller 103 across all of the baby-busses 107, each one of which may consist of

sixteen data lines and a strobe line 112 (to be discussed further below). The controller 103 also uses a single command bus 111 which connects the controller 103 with all memory chips 105. Typically, the command bus 111 will
5 include each of an address bus 113, a system clock 115, a data read signal 117, a data write signal 119, and other lines as well (not shown) such as various interrupt lines, chip select circuitry and the like. Preferably, each memory chip 105 is a random access memory and,
10 specifically, a dynamic random access memory, requiring periodic refresh. Since the memory devices can be affected by temperature, each device 105 is preferably configured both upon system power-up and also at periodic intervals for the purposes of synchronizing memory retrieval.

15 Each chip 105 includes an input buffer (not seen in FIG. 3) which is loaded with a number after the controller 103 has polled all of the memory chips. The number is derived from a number of elapsed system clock cycles, counted by the controller, corresponding to the slowest
20 chip in the system 101 to respond. Preferably, the controller 103 during the configuration mode separately enables each memory chip 105, by programming a configuration command into a mode register of the chip (seen in FIG. 5, and described below), which then causes
25 only that chip to be active and to enter a mode where it listens to its corresponding baby-bus. When a chip 105 in this mode detects an input strobe (including a predetermined data transmission to it along the data bus), it responsively transmits a predetermined response word
30 back onto the data bus, together with the output strobe 112, which is then used to latch the predetermined response word into the controller 103. Receipt of the output strobe 112 and the predetermined response word indicates to a comparison mechanism 118 of the controller 103 that it

should freeze an internal timer, which has been incremented with each system clock pulse elapsed since the controller issued the input strobe, and load that counter's contents into memory internal to the controller, both into a memory spot 114 representing a maximum count (if the counter represents a maximum), and also into a slot 116 dedicated to the particular chip 105. Once the controller has polled all of the memory chips 105, it then calculates an offset for each chip by which each chip is to delay its output, such that data read operations provide data to the memory controller at the maximum response time. In other words, the controller 103 subtracts the response time for each chip from the maximum, to obtain the offset, and it then programs each chip 105 with the appropriate offset.

FIG. 4 is a block diagram a single chip 105 from FIG. 3. In particular, the preferred chip is a DRAM 123 which includes an internal delay mechanism 121 for synchronizing memory response with other chips; preferably, this delay mechanism includes an output latch and an internal counter and the internal buffer. The DRAM 123 is coupled to the memory controller 103 via a number of communication paths, including (a) the command bus 111, (b) a 16-bit baby-bus 107, (c) the strobe signal 117, (d) a command strobe 125, (e) the system clock 115 and (f) a chip select signal, not seen in FIG. 4. When the system clock 115 is received by the DRAM 123, an internal phase locking mechanism 127 is used to provide a locally-derived clock signal 129 for timing internal operations of the DRAM and for defining a local time zero for coordinating memory operations.

When it is desired to configure a particular chip, e.g., the DRAM 123, during a configuration mode, the controller 103 enables the specific DRAM, simultaneously

disabling all other memory chips along the same baby-bus 107. This enablement is achieved by causing the controller 103 to send a configuration command to the DRAM via the command bus 111, which together with a command strobe, 5 causes the DRAM to latch the command into internal timing and control logic 131. Since the configuration command, as indicated by FIG. 4, must pass through a second latch 133 before it reaches the internal timing and control logic 131, its arrival will be synchronized with the local clock 10 signal 129 of the DRAM 123. The configuration command, for example, may be indicated by the most significant bits only of the address bus, and it does not require transmission of data along the baby-bus 107. These most significant bits are then coupled by the internal timing and control logic 15 131 into a mode register 132 and used to direct the DRAM 123 to enter the configuration mode.

In the configuration mode, the DRAM will simply monitor the baby-bus 107 for a predetermined data word, which will be accompanied by an input strobe to strobe the 20 data word into a sync detect register 137. In response to detection of the data word, the timing and control logic will immediately cause a predetermined response word to be fed to an output latch 139 from a memory space 141, which the timing and control logic will enable via a sync control 25 line 143. The predetermined response word will thereby be immediately gated back onto the baby-bus 107 and sent to the controller, usually on the local clock pulse immediately subsequent to presentation of the response word to the output latch 139. Accompanied by an output strobe, 30 the predetermined response word is latched into the controller immediately upon receipt, and used to stop the internal timer of the controller's comparison mechanism 118.

FIG. 5 is used to indicate the results of the controller's determination of a maximum delay time across all of the chips 105 of FIG. 3. In particular, as was the case in connection with FIG. 2, it is desired for the controller 103 of FIG. 3 to perform data reads on successive clock cycles. In this regard, the controller sends data reads "X" and "Y" (designated by the reference numerals 145 and 147 in FIG. 5B) to two different memory devices which are serially coupled to the controller. These devices may consist of single chips 105 along a single baby-bus 107-1, or alternatively, may consist of plural chips coupled to the controller via different baby-busses 107. The data reads 145 and 147 are the first and third reads of a train of three consecutive data reads, with the middle data read indicated by a dash "-" in FIG. 5B. As with FIG. 2, each data read operation has a propagation time associated with it, and the operations arrive at a corresponding memory device in accordance with the corresponding propagation time. However, implementing the principles of the present invention, the memory device associated with command "Y" is programmed with an offset 149 representing a two clock cycle delay. Thus, that memory device will use its local clock to impose a delay of two clock cycles to the output latch (139 in FIG. 4). The offset corrects each individual chip to respond in accordance with the maximum response time for all of the chips, e.g., each chip has its response slowed to be in sync with the slowest chip. Consequently, presuming that all three of the memory devices are configured in accordance with the present invention, each one of the three data reads will be received on successive clock cycles, as indicated by the reference numeral 151 in FIG. 5E.

Returning to FIG. 4, the controller programs delay by loading an offset into the DRAM 123; it performs this programming by writing a number of clock pulses into an internal buffer 153 of the DRAM. Subsequently, when the
5 DRAM 123 receives a data read command outside of the configuration mode, it processes the command in the normal fashion. However, when the product of the read command is presented at the output latch, that product is not immediately gated onto the data bus on the subsequent local
10 clock pulse. Rather, the internal timing and control logic 121 withholds enablement of the sync control line 143 to delay output by an amount indicated by the buffer. Preferably, the buffer is combined with a timer for this purpose, which begins a countdown upon the detection of
15 each read command received by the chip.

Importantly, there are a number equivalent implementations of this delay structure that will readily occur to those of skill in the art; for example, offset could be achieved by designing each memory device and the
20 system to provide output at a predetermined clock cycle, but to program a number to make the output available "early" in response to buffer contents. Alternatively, each memory device could have multiple modes, with delay achieved by selecting one of several modes. Alternatively,
25 delay could be implemented by circuit external to each memory device. These environments, as well as other modifications that will occur to those having skill in memory system design, are contemplated as being within the spirit of the present invention.

30 After the DRAM 123 has been programmed with the offset, the controller preferably calculates a correction which is used to refine the offset. Specifically, in the system of FIG. 3, each data buss (each baby-bus 107) and

the corresponding command bus 111 may have different path lengths. Since the offset was originally calculated using polling and reply both over the data bus only (the baby-bus 107), and since read commands are normally issued along the command bus, the offset may not correctly reflect desired delay. Consequently, the correction for each chip is used to refine the delay time associated with each chip. 105 to correct this discrepancy, and the correction may differ from chip-to-chip.

10 To calculate the correction, the controller 103 sends a follow-up write command to each DRAM; the write command includes both the actual command transmitted along the command bus 111, as well as data transmitted along the corresponding baby-bus 107. With reference to FIG. 4, as
15 the command is received by the DRAM 123 from the command, it is provided to a first sync detect circuit 155, which provides an indication of a received command to a command state counter 157. The command state counter 157 then begins counting with each local clock pulse. The data
20 corresponding to the command received from the corresponding baby-bus 107 is also fed to a second sync detect circuit 159, which similarly provides an indication of received data to a data state counter 161. This data state counter 161 also begins counting with each local
25 clock pulse. Thus, receipt of either the write command, or its corresponding write data, will trigger one or the other of the counters 157 or 161 to begin counting. Receipt of both indications from the sync detect circuits 155 and 159 is used to gate, via a subtract circuit 163, a difference
30 in number of clock cycles between command receipt, and an output reflecting this difference is provided to the timing and control logic 131 of the DRAM 123. This difference represents the correction to the offset in number of clock cycles. Thus, depending upon whether the command bus is

slower or faster than the data bus, in terms of propagation time, the timing and control logic 131 will modify the contents of the internal buffer 153 accordingly. Consequently, once the offset and correction are

5 calculated, the contents of the internal buffer 153 will thereafter precisely reflect a desired number of offset clock cycles for achieving synchronization across multiple memory devices.

What has been described is a memory system and

10 device for achieving synchronous data reply, at the controller, from many memory devices, such as from DRAM chips. Using the preferred embodiment described above, one may achieve substantially synchronous memory response either from parallel memory devices, e.g., as connected

15 across multiple baby-busses 107, and providing bit groups of different significance for one data word fetch (e.g., for a 64-bit word fetch), or for avoiding bus contention between multiple memory devices coupled to the same data bus, as explained with reference to FIGS. 2 and 4.

20 Notably, the system described herein performs correction for only integer clock cycles, and does not achieve phase correction between different memory chips. The system described herein, however, is compatible with systems for achieving phase synchronization, such as described in U.S.

25 Patent No. 4,998,262, which is hereby incorporated by reference.

Having thus described an exemplary embodiment of the invention, it will be apparent that further alterations, modifications, and improvements will also

30 occur to those skilled in the art. Such alterations, modifications, and improvements, though not expressly described or mentioned above, are nonetheless intended and implied to be within the spirit and scope of the invention.

Accordingly, the foregoing discussion is intended to be illustrative only; the invention is limited and defined only by the various following claims and equivalents thereto.

CLAIMS

1 1. In a digital memory system (101) having a
2 controller (103) and at least two memory devices (105),
3 with at least one data bus (107) coupling the controller
4 (103) with each memory device (105), the controller (103)
5 also having a (115) by which the controller (103) times its
6 operations, an improvement comprising:

7 a buffer (153) in each memory device (105), each
8 buffer (153) adapted to store a number (149)
9 corresponding to an amount of time for which the
10 corresponding buffer (153) is to delay data output;

11 a routine run by the controller (103) which causes
12 the controller

13 to, for each memory device,

14 prompt the memory device for a response,
15 and

16 count a number of clock pulses at the
17 controller (103) between prompting of the
18 memory device (105) and receipt of the
19 corresponding response from the memory
20 device (105),

21 to determine a maximum delay measured by the
22 maximum number of clock pulses needed for
23 response across the at least two memory devices
24 (105), and

25 to cause the buffer (153) of each memory device
26 (105) to store a number sufficient to cause
27 data from the corresponding memory device (105)

28 to be received at the controller (103) at the
29 maximum delay time; and

30 circuitry (121) in each memory device (105) that
31 delays data output in response to the number stored
32 in the corresponding buffer (153).

1 2. An improvement according to claim 1, wherein each
2 memory device (105) receives the system clock (115) and
3 generates a local clock (129) in response thereto, said
4 improvement further comprising:

5 an output strobe (112) issued by each memory device
6 (105) in response to being prompted, the output
7 strobe (112) being locally derived from the system
8 clock (115) at the location of the corresponding
9 memory device (105) and coupled to the controller
10 (103) so as to strobe arrival of the response from
11 the particular memory device (105).

1 3. An improvement according to claims 1 or 2, wherein
2 the system bus includes a data bus (107) and a command bus
3 (111) having different path lengths, said improvement
4 further comprising:

5 a delay configuration mode of each memory device
6 (105), the delay configuration mode being
7 instructed through the command bus (111) and
8 causing the corresponding memory device (105) to

9 wait for presence of a data signal on the data
10 bus (107) and an input strobe (125) to accept
11 the data, the data signal forming at least part
12 of the prompt from the controller (103), and

13 issue the response to the controller (103) in
14 reply to the data signal.

1 4. An improvement according to claim 3, further
2 comprising:

3 a delay adjustment mechanism within each memory
4 device, the delay adjustment mechanism receiving a
5 command from the command bus (111) and data from
6 the data bus (107), and calculating a deviation in
7 time to the number for the corresponding memory
8 device (105), the number correcting for difference
9 in travel time between the command bus (111) and
10 the data bus (107).

1 5. An improvement according to claim 1, 2, 3 or 4,
2 said improvement further comprising:

3 arrangement of at least two memory devices (105) on
4 a single, shared data bus (107).

1 6. An improvement according to claim 1, 2, 3, 4 or 5,
2 said improvement further comprising:

3 arrangement of at least two memory devices (105) on
 different, parallel data busses (107).

1 7. A method of synchronizing response in a memory
2 system (101) which includes a plurality of memory devices
3 (105) coupled to a controller (103) via at least one data
4 bus (107), wherein the controller (103) performs data read
5 operations from the memory devices (105) according to its
6 own system clock (115), and the data read operations have
7 an associated travel time between the controller (103) and
8 corresponding memory devices (105), said method comprising:

9 individually enabling each memory device (105)
10 along each data bus (107);

11 while a particular memory device (105) is enabled,
12 sending a signal to the particular memory device
13 (105) from the controller (103);

14 for each memory device (105), sending a reply to
15 the controller (103) in response to the
16 corresponding signal;

17 counting an elapsed time between sending the signal
18 and receipt of the reply to obtain a response time,
19 for each memory device (105);

20 in response to the response times for the memory
21 devices (105), determining a maximum response time;
22 and

23 causing each of the memory devices (105) to provide
24 a data reads to the controller (103) at a
25 predetermined time from command, selected in
dependence upon the maximum response time.

1 8. A method according to claim 7, wherein each memory
2 device (105) includes a buffer (153) and a delay mechanism
3 (121) causing the memory device (105) to delay its output
4 by an amount determined in response to contents of the
5 buffer (153), wherein:

6 causing each of the memory devices (105) includes
7 loading each buffer (153) with a corresponding
8 number; and

9 delaying with the delay mechanism (121) an output
10 of each memory device (105) in response to the
11 contents of the corresponding buffer (153).

1 9. A method according to claims 7 or 8, wherein at
2 least one memory device (105) accepts data write commands
3 from the controller (103) and includes a mode register
4 (132) that enables the at least one memory device (105) and
5 causes it to await a polling command from the controller
6 (103) along at least one data bus (107) associated with the
7 at least one memory device (105), and the memory system
8 (101) further includes at least one command bus (111)
9 coupling the at least one memory device (105) to the
10 controller (103), said method further comprising:

11 using the command bus (111) to command the mode
12 register (132) to cause the at least one memory
13 device (105) to await the polling command;

14 sending the polling command to the at least one
15 memory device (105), the system controller (103),
16 in association with sending the polling command,
17 counting system clock cycles;

18 detecting the polling command at the memory device
19 (105) and responsively sending the reply to the
20 controller (103);

21 detecting the reply at the controller (103) and
22 identifying a number of elapsed clock cycles; and

23 computing the number for the at least one memory
24 device (105) therefrom.

1 10. A method according to claims 7, 8, or 9, further
2 using a system clock signal (115) issued via the controller
3 (103), wherein the at least one memory device (105) is
4 coupled to the system clock signal (115) and provides an
5 internal clock (129) based on the system clock signal
6 (115), and further, provides a strobe signal (112) as an
7 output, said method further comprising:

8 issuing the strobe signal (112) in synchronization
9 with issuing the reply; and

10 using the strobe signal (112) to indicate to the
11 system controller (103) that the reply is arriving
12 at the system controller (103).

1 11. A method according to claims 7, 8, 9, or 10,
2 wherein each memory device (105) further includes a
3 predetermined data word stored internally, said method
4 further comprising:

5 for each memory device (105), in response to the
6 signal, coupling the predetermined data word to at
7 least one associated data bus (107).

1 12. A method according to claim 7, 8, 9, 10 or 11, said
2 method further comprising:

3 arranging at least two memory devices (105) on a
4 single, shared data bus (107).

1 13. A method according to claim 7, 8, 9, 10, 11 or 12,
2 said method further comprising:

3 arranging at least two memory devices (105) on
4 different, parallel data busses (107).

1 14. A memory device (105) for use in a memory
2 system (101) having a memory controller (103) and a system
3 bus, wherein the memory controller (103) polls each of
4 several memory elements (105) to determine an associated
5 command response time and determines a desired response
6 time, comprising:

7 a memory space (141) wherein data is stored;

8 an output path for coupling data from the memory
9 space to the system bus in response to a data read
10 command;

11 memory space drivers that accesses data locations
12 within the memory space (141) and that couple
13 contents of selected ones of the data locations to
14 the output path in response to a data read command;

15 a clock mechanism (127) that provides a clock
16 signal (129) to synchronize internal operations of
17 said memory device (105);

18 a buffer (153) that stores a number of clock
19 pulses; and

20 a delay mechanism (121) coupled to the buffer
21 (153), the delay mechanism (121) causing delay of
22 at least one of said output path and said memory
23 space drivers by an amount corresponding the number
24 of clock pulses;

25 wherein said device is adapted for synchronization
26 with other memory elements (105), in response to
27 the controller (103) causing the buffer (153) to be
28 loaded with a number that causes equalization of
29 command response time for said device to the
30 desired response time.

1 15. A memory device according to claim 14, wherein the
2 memory system (101) further includes a data bus (107) and a
3 system clock (115) issued by the system controller (103),
4 and wherein:

5 said device further comprises

6 a mode register (132) that is selectively
7 programmed by command from the system
8 controller (103) to cause said device to enter
9 a calibration mode,

10 a predetermined response word stored within
11 said device, the predetermined response word
12 being coupled to the data bus (107) in response
13 to a polling command from the system controller
14 (103), and

15 an output strobe (112) of said device, the
16 output strobe (112) being locally derived at
17 said device from the system clock (115); and

18 the calibration mode causes said device to await
19 the polling command and, in response to detection
20 of the polling command by said device on the data
21 bus (107), to couple the predetermined response
22 word and the echo strobe onto the data bus (107).

1 16. A memory device according to claims 14 or 15,
2 wherein the system bus further includes a command bus
3 (111), and wherein:

4 said device further receives the number from the
5 system controller (103) and stores it in the buffer
6 (153); and

7 said device further comprises
8 a comparison mechanism (163) for comparing time
9 of receipt of at least one signal along the
10 command bus with time of receipt of a
11 corresponding signal along the data bus, and
12 for providing an output, and

13 an adder to which the output is coupled, the
14 adder (131) modifying the number by the output
15 to compensate for difference between travel
16 time along the data bus and travel time along
17 the command bus.

1 17. A memory device according to claims 14, 15, or 16,
2 wherein the memory device is a random access memory.

1 18. A memory device according to claim 14, 15, or 16,
2 wherein the memory device is a dynamic random access
3 memory.

1 19. A memory system, comprising:

2 a memory controller (103);

3 a first memory device (105) having a buffer (153),
4 and internal delay means (121) for delaying output
5 in response to contents of the buffer (153);

6 a second memory device;

7 a system bus coupling one of commands and data
8 between the memory controller (103) with each of
9 the first memory device and second memory device,
10 the system bus having difference in travel times
11 between the memory controller with the first memory

12 device, and the memory controller with the second
13 memory device;

14 a calibration mode of the controller (103), which
15 causes the controller (103) to poll each of the
16 first memory device and the second memory device
17 and to time response from each of the first memory
18 device and the second memory device; and

19 a comparison device of the controller (118), which
20 compares during the calibration mode the timed
21 response from each of the first memory device (105)
22 and second memory device and calculates a number
23 corresponding thereto;

24 wherein the calibration mode of the controller
25 (103) causes the controller to store the number in
26 the buffer (153), such that the internal delay
27 means (121) delays output in response to the
28 number.

1 20. A memory system according to claim 19 wherein:
2 both of the first and second memory devices are
3 coupled to the controller (103) on a single, shared
4 data bus (107).

1 21. A memory system according to claims 19 or 20,
2 wherein:
3 both of the first and second memory devices are
4 coupled to the controller (103) on different,
5 parallel data busses (107).

1 22. A memory system according to claims 19 or 20
2 wherein:

3 both of the first and second memory devices are
4 coupled to the controller (103) on a single, shared
5 data bus (107);



6 said memory system (101) further comprises a third
7 memory device which is coupled to the controller
8 (103) on a second data bus, different from the
9 shared data bus, the third memory also having a
10 buffer (153);

11 the calibration mode of the controller (103) causes
12 to poll the third memory device;

13 the comparison device (118) additionally compares a
14 timed response from the third memory device with
15 time response for at least one of the first and
16 second memory devices;

17 the calibration mode of the controller (103) causes
18 the controller to store a second number in the
19 buffer (153) of the third memory, such that the
20 third memory may be synchronized for simultaneous
21 data delivery at the controller with one of the
22 first and second memory devices.

FIG. 1
(Prior Art)






FIG. 4

INTERNATIONAL SEARCH REPORT

Int'l Application No

PCT/US 97/18128

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F13/16

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 339 224 A (I. B. M.) 2 November 1989 see page 3, line 15 - page 4, line 17 see page 5, line 19 - page 6, line 43; claims; figures 2,4 ---	1,7,14, 19
A	US 5 509 138 A (CHRISTINE G. CASH) 16 April 1996 see column 2, line 26 - line 64 see column 8, line 65 - column 9, line 22; claims; figure 4 ---	1,7
A	EP 0 238 090 A (NEC CORPORATION) 23 September 1987 see column 2, line 25 - column 3, line 7 see column 9, line 51 - column 11, line 57; claims; figures -----	1,7



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

18 February 1998

Date of mailing of the international search report

25/02/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Soler, J

INTERNATIONAL SEARCH REPORT

Information on patent family members

Int. :ional Application No

PCT/US 97/18128

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 339224 A	02-11-89	JP 2012541 A US 5301278 A	17-01-90 05-04-94
US 5509138 A	16-04-96	NONE	
EP 238090 A	23-09-87	DE 3752017 D DE 3752017 T JP 63012049 A US 4870562 A	20-03-97 28-08-97 19-01-88 26-09-89